PCT/EP2004/052502

The Patent Office

PATENTS • DESIGNS • COPYRIGHT • TRADE MARKS

16 NOV 2004

INVESTOR IN PEOPLE

REC'D 13 JAN 2005	
WIPO	PCT

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

PRIORITY DOCUMENT
SUBMITTED OR TRANSMITTED IN COMPLIANCE WITH RULE 17.1 (a) OR (b)

Signed [signature]

Dated 22 October 2004

BEST AVAILABLE COPY

Patents Form 1/77

Patents Act 1977
(Rule 16)

Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)

9 OCT 2003

The Patent Office

Cardiff Road
Newport
South Wales
NP10 8QQ

1.	Your reference	P690GB
2.	Patent application number *(The Patent Office will fill this part in)*	0323662.7
3.	Full name, address and postcode of the or of each applicant *(underline all surnames)*	UNIVERSITY OF YORK HESLINGTON YORK YO1 5DD GB
	Patents ADP number *(if you know it)*	04169546006
	If the applicant is a corporate body, give the country/state of its incorporation	United Kingdom
4.	Title of the invention	IMAGE RECOGNITION
5.	Name of your agent *(if you have one)*	STANLEYS
	"Address for service" in the United Kingdom to which all correspondence should be sent *(including the postcode)*	Kings Court 12 King Street LEEDS LS1 2HL
	Patents ADP number *(if you know it)*	07902190001 ✓

		Country	Priority application number *(if you know it)*	Date of filing *(day / month / year)*
6.	Priority: Complete this section if you are declaring priority from one or more earlier patent applications, filed in the last 12 months.			

		Number of earlier UK application	Date of filing *(day / month / year)*
7.	Divisionals, etc: Complete this section only if this application is a divisional application or resulted from an entitlement dispute *(see note f)*		

8. Is a Patents Form 7/77 (Statement of inventorship and of right to grant of a patent) required in support of this request? — YES

Answer YES if:
a) any applicant named in part 3 is not an inventor, or
b) there is an inventor who is not named as an applicant, or
c) any named applicant is a corporate body.
Otherwise answer NO (See note d)

ents Form 1/77

9. Accompanying documents: A patent application must include a description of the invention. Not counting duplicates, please enter the number of pages of each item accompanying this form:

Continuation sheets of this form	
Description	16
Claim(s)	
Abstract	
Drawing(s)	6

10. If you are also filing any of the following, state how many against each item.

Priority documents	
Translations of priority documents	
Statement of inventorship and right to grant of a patent (Patents Form 7/77)	
Request for a preliminary examination and search (Patents Form 9/77)	
Request for a substantive examination (Patents Form 10/77)	
Any other documents (please specify)	

11. I/We request the grant of a patent on the basis of this application.

Signature(s) STANLEYS Date 9 October 2003

12. Name, daytime telephone number and e-mail address, if any, of person to contact in the United Kingdom — David Stanley 01481 824411

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

a) If you need help to fill in this form or you have any questions, please contact the Patent Office on 08459 500505.

b) Write your answers in capital letters using black ink or you may type them.

c) If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.

d) If you have answered YES in part 8, a Patents Form 7/77 will need to be filed.

e) Once you have filled in the form you must remember to sign and date it.

f) Part 7 should only be completed when a divisional application is being made under section 15(4), or when an application is being made under section 8(3), 12(6) or 37(4) following an entitlement dispute. By completing part 7 you are requesting that this application takes the same filing date as an earlier UK application. If you want the new application to have the same priority date(s) as the earlier UK application, you should also complete part 6 with the priority details.

IMAGE RECOGNITION

This invention relates to the recognition of images, and is concerned, particularly although not exclusively, with the recognition of natural images.

By "natural image" is meant an image of an object that occurs naturally – for example, an optical image such as a photograph, as well as images of other wavelengths – such as x-ray and infra-red, by way of example. The natural image may be recorded and/or subsequently processed by digital means, but is in contrast to an image – or image data – that is generated or synthesised by computer or other artificial means.

The recognition of natural images can be desirable for many reasons. For example, distinctive landscapes and buildings can be recognised, to assist in the identification of geographical locations. The recognition of human faces can be useful for identification and security purposes. The recognition of valuable animals such as racehorses may be very useful for identification purposes.

In this specification, we present in preferred embodiments of the invention a new approach to face recognition using a variety of three-dimensional facial surface representations generated from a University of York (UofY) /Cybula 3D Face Database. By applying principal component analysis to three-dimensional surface structure, we show that high levels of accuracy can be achieved when performing recognition on a large database of 3D face models, captured under conditions that present typical difficulties to the more conventional two-dimensional approaches. Results are presented as false acceptance rates and false rejection rates, taking the equal error rate as a single comparative value. We identify the most effective surface

representations and distance metrics to be used in such application areas as security, surveillance, data compression and archive searching.

Despite significant advances in face recognition technology, it has yet to achieve the levels of accuracy required for many commercial and industrial applications. Although some face recognition systems proclaim extremely low error rates in the test environment, these figures often increase when exposed to a real world scenario. The reasons for these high error rates stem from a number of well-known sub-problems that have never been fully solved. Face recognition systems are highly sensitive to the environmental circumstances under which images are captured. Variation in lighting conditions, facial expression and orientation can all significantly increase error rates, making it necessary to maintain consistent image capture conditions between query and gallery images for the system to function adequately. However, this approach eliminates some of the key advantages offered by face recognition: a passive biometric in the sense that it does not require subject co-operation.

In preferred embodiments, we use 3D face models that eliminate some of the problems commonly associated with face recognition. By relying purely on geometric shape, rather than the colour and texture information available in two-dimensional images, we render the system invariant to lighting conditions, at the expense of loosing the distinguishing features only available in colour and texture data. In addition, the ability to rotate a facial structure in three-dimensional space allows for compensation of variations in pose, aiding those methods requiring alignment prior to recognition.

Here we use facial surface data for the first time, taken from 3D face models, as a substitute for the more familiar two-dimensional images. We take a well-known method of face recognition, namely the eigenface approach

described by Turk and Pentland [1, 2] and adapt it for use on the new three-dimensional data. We identify the most effective methods of recognising faces using three-dimensional surface structure.

In order to test this method of face recognition, we have used a large database of 3D face models. However, until recently, methods of 3D model generation have usually employed the use of laser scanning equipment. Such systems (although highly accurate) are often slow, requiring the subject to remain perfectly still. Stereo vision techniques are able to capture at a faster rate without using lasers, but feature correlation requires regions of contrast and stable local texture; something that cheeks and forehead distinctly lack. For these reasons, three-dimensional face recognition has remained relatively unexplored, when compared to the wealth of research focusing on two-dimensional face recognition. Although some investigations have experimented with 3D data [3, 4, 5, 6], they have had to rely on small tests sets of 3D face models or used generic face models to enhance two-dimensional images prior to recognition [7, 8, 9]. However, this research demonstrates that the use of three-dimensional information has the potential to improve face recognition well beyond the current state of the art. With the emergence of new three-dimensional capture equipment, population of a large 3D face database has now become viable and being undertaken at the UofY/Cybula as part of a project facilitating research into three-dimensional face recognition technology.

Previous research has explored the possibilities offered by three-dimensional geometric structure to perform face recognition. To date, the research has focused on two-dimensional images, although some have attempted to use a-priori knowledge of facial structure to enhance these existing two-dimensional approaches. For example, Zhao and Chellappa [7] use a generic 3D face model to normalise facial orientation and lighting direction in two-

dimensional images. Using estimations of light source direction and pose, the 3D face model is aligned with the two-dimensional face image and used to project a prototype image of the frontal pose equivalent, prior to recognition by linear discriminant analysis. Though this approach, recognition accuracy on the test set is increased from approximately 81% (correct match within rank of 25) to 100%. Similar results are witnessed in the Face Recognition Vendor Test [10], showing that pose correction using Romdhani, Blanz and Vetter's 3D morphable model technique [8] reduces error rates when applied to the FERET database [11].

Blanz, Romdhani and Vetter [9] take a comparable approach, using a 3D morphable face model to aid in identification of 2D face images. Beginning with an initial estimate of lighting direction and face shape, Romdhani et al iteratively alters shape and texture parameters of the morphable face model, minimising difference to the two-dimensional image. These parameters are then taken as features for identification.

Although the methods discussed show that knowledge of three-dimensional face shape can improve two-dimensional face recognition systems by improving normalisation, none of the methods mentioned so far use actual geometric structure to perform recognition. Whereas Beumier and Acheroy [3] make direct use of such information, generating 3D face models using an approach based on structured light deformation. Beumier and Acheroy test various methods of matching 3D face models; few of which were successful. Curvature analysis proved ineffective, and feature extraction was not robust enough to provide accurate recognition. However, Beumier and Acheroy were able to achieve reasonable error rates using curvature values of vertical surface profiles. Verification tests carried out on a database of 30 people produced equal

error rates between 7.25% and 9% on the automatically aligned surfaces and between 6.25% and 9.5% when manual alignment was used.

Chua et al [6] take a different approach, applying non-rigid surface recognition techniques to the face structure. An attempt is made to identify and extract rigid areas of facial surfaces, creating a system invariant to facial expression. The characteristic used to identify these rigid areas and ultimately distinguish between faces is the point signature, which describes depth values surrounding local regions of specific points on the facial surface. The similarity of two face models is computed by identifying and comparing a set of unique point signatures for each face. Identification tests show that the probe image is identified correctly for all people when applied to a test set of 30 depth maps of 6 different people.

Coombes et al [12] investigate a method based on differential geometry. Curvature analysis is applied to a depth map of the facial surface; segmenting the surface into one of eight fundamental types: peak, ridge, saddle ridge, minimal, pit, valley, saddle valley and flat. Coombes et al suggest that two faces may be distinguished by comparing which curve types classification of correlating regions. A quantitative analysis of the average male and female face structure shows distinct differences in chin, nose, forehead shape and cheek bone position between faces of different gender.

Another method, proposed by Gordon [5], incorporates feature localisation. Using both depth and curvature information extracted from three-dimensional face models, Gordon identifies a number of facial features, from which a set of measurements are taken, including head width, numerous nose dimensions and curvatures, distance between the eyes and eye width. These features are evaluated using fisher's linear discriminant, determining the

discriminating ability of each individual feature. Gordon's findings show that the head width and nose location are particularly important features for recognition, whereas eye widths and nose curvatures are less useful. Recognition is performed by means of a simple euclidean distance measure in feature space. Several combinations of features are tested using a database of 24 facial surfaces taken from 8 different people, producing results ranging from 70.8% to 100% correct matches.

According to one aspect of the present invention, there is provided a method of recognising a natural image, comprising the steps of generating a depth map of the image, generating eigenvectors and eigenvalues from the depth map, and recognising the image by those eigenvectors and eigenvalues.

In another aspect, the invention provides a device for recognising a natural image, comprising means for generating a depth map of the image, means for generating eigenvectors and eigenvalues from the depth map, and means for recognising the image by those eigenvectors and eigenvalues.

A method or device as above may further include any or all of the features or method steps disclosed in this specification (including the drawings), which may be combined in any combination, except combinations where at least some of such features and/or steps are mutually exclusive.

For a better understanding of the invention, and to show how embodiments of the same may be carried into effect, reference will now be made, by way of example, to the accompanying diagrammatic drawings, in which:

Figure 1 shows examples of face models taken from a 3D face database;

Figure 2 shows orientation of a raw 3D face model (left) to a frontal pose (middle) and facial surface depth map (right);

Figure 3 shows an average depth map (left most) and first eight eigensurfaces;

Figure 4 is a graph showing false acceptance rate and false rejection rate for baseline 3D face recognition systems using facial surface depth maps and a range of distance metrics;

Figure 5 is a diagram of verification test procedure;

Figure 6 is a graph showing false acceptance rate and false rejection rate for 3D face recognition systems using optimum surface representations and distance metrics;

Figure 7 is a chart to show Equal error rates of 3D face recognition systems using a variety of surface representations and distance metrics; and

Figure 8 shows brief descriptions of surface representations with convolution kernels used.

As mentioned previously, there is little three-dimensional face data publicly available at present and nothing towards the magnitude of data required for development and testing of three-dimensional face recognition systems. Therefore, we have collected a new database of 3D face models, collected at UofY/Cybula as part of an ongoing project to provide a publicly available 3D Face Database of over 1000 people for face recognition research. The 3D face models are generated using a stereo vision technique enhanced by light projection to provide a higher density of features. Each face model requires a

single shot taken with a 3D camera, from which the model is generated in sub-second processing time.

For the purpose of these evaluation, we use a subset of the 3D face database, acquired during preliminary data acquisition sessions. This set consists of 330 face models taken from 100 different people under the conditions shown in Fig. 1.

During capture, no effort was made to control lighting conditions. In order to generate face models at various head orientations, subjects were asked to face reference points positioned roughly 45 above and below the camera, but no effort was made to enforce a precise angle of orientation. Examples of the face models generated for each person are shown in Fig. 1.

3D face models are stored in the OBJ file format (a common representation of 3D data) and orientated to face directly forwards using our orientation normalisation algorithm (not described here) before being converted into depth maps. The database is then separated into two disjoint sets: the training set consisting of 40 depth maps of type 01 (see Fig. 1) and a test set of the remaining 290 depth maps, consisting of all capture conditions shown in Fig. 1. Both the training set and test set contain subjects of various race, age and gender and nobody is present in both the training and test sets.

In previous work we have shown that the use of image processing techniques can significantly reduce error rates of two-dimensional face recognition methods, by removing unwanted features caused by environmental capture conditions. Much of this environmental influence is not present in the 3D face models, but pre-processing may still aid recognition by making distinguishing features more explicit. In this section we describe a number of

surface representations, which may affect recognition error rates. These surfaces are derived by pre-processing of depth maps, prior to both training and test procedures, as shown in Fig. 4.

In our approach we define a '3D surface space' by application of principal component analysis to the training set of facial surfaces, taking a similar approach to that described by Turk and Pentland [1] and used in previous investigations.

Consider our training set of facial surfaces, stored as orientation normalised 60x105 depth maps. Each of these depth maps can be represented as a vector of 6300 elements, describing a single point within the 6300 dimensional space of all possible depth maps. What's more, faces with a similar geometric structure should occupy points in a comparatively localised region of this high dimensional space. Continuing this idea, we assume that different depth maps of the same face project to nearby points in space and depth maps of different faces project to far apart points. Ideally, we wish to extract the region of this space that contains facial surfaces, reduce the dimensionality to a practical value, while maximising the spread of facial surfaces within the depth map subspace.

In order to define a space with the properties mentioned above, we apply principal component analysis to the training set of M depth maps (in our case M = 40) $\{\Gamma_1, \Gamma_2, \Gamma_3, \ldots \Gamma_M\}$, computing the covariance matrix,

$$C = \frac{1}{M}\sum_{n=1}^{M} \Phi_n \Phi_n^T = AA^T \qquad A = [\Phi_1 \Phi_2 \Phi_3 \ldots \Phi_M] \quad \Phi_n = \Gamma_n - \Psi \quad \Psi = \frac{1}{M}\sum_{n=1}^{M}\Gamma_n \qquad (1)$$

Where Φ_n is the difference of the nth depth map from the average ψ. Eigenvectors and eigenvalues of the covariance matrix are calculated using standard linear methods. The resultant eigenvectors describe a set of axes within the depth map space, along which most variance occurs within the training set and the corresponding eigenvalues represent the degree of this variance along each axis. The M eigenvectors are sorted in order of descending eigenvalues and the $M`$ greatest eigenvectors (in our system $M` = 40$) are chosen to represent surface space. The effect is that we have reduced the dimensionality of the space to M`, yet maintained a high level of variance between facial surfaces throughout the depth map subspace.

We term each eigenvector an eigensurface, containing 6300 elements (the number of depth values in the original depth maps) which can be displayed as range images of the facial surface principal components, shown in Fig. 3.

Once surface space has been defined we project any face into surface space by a simple matrix multiplication using the eigenvectors calculated from the covariance matrix in equation 1:

$$\omega_k = u_k^T(\Gamma - \Psi) \qquad \textit{for } k = 1 \ldots M`. \tag{2}$$

where u_k is the kth eigenvector and ω_k is the kth weight in the vector $\Omega^T = [\omega_1, \omega_2, \omega_3, \ldots \omega_{M`}]$. The M` coefficients represent the contribution of each respective eigensurface to the projected depth map. The vector Ω is taken as the 'face-key' representing a person's facial structure in surface space and compared by either euclidean or cosine distance metrics.

$$d_{euclidean} = \|\Omega_a - \Omega_b\| \qquad d_{cosine} = 1 - \frac{\Omega_a^T \Omega_b}{\|\Omega_a\|\|\Omega_b\|} \quad (3)$$

In addition, we can also divide each face-key by its respective eigenvalues, prior to distance calculation, removing any inherent dimensional bias and introducing two supplementary metrics, the mahalanobis distance and weighted cosine distance. An acceptance (the two facial surfaces match) or rejection (the two surfaces do not match) is determined by applying a threshold to the calculated distance. Any comparison producing a distance below the threshold is considered an acceptance. In order to evaluate the effectiveness of the face recognition methods, we carry out 41,905 verification operations on the test set of 290 facial surfaces, computing the error rates produced (see Fig. 4). Each surface in the test set is compared with every other surface, no image is compared with itself and each pair is compared only once (the relationship is symmetric).

False acceptance rates and false rejection rates are calculated as the percentage of incorrect acceptances and incorrect rejections after applying the threshold. Applying a range of thresholds produces a series of FAR, FRR pairs, which are plotted on a graph as shown for our benchmark system in Fig. 5. The equal error rate can be seen as the point where FAR equals FRR.

We now present the results gathered from testing the three-dimensional face recognition methods on the test set of 290 facial surfaces. The results are presented by error curves of FAR vs. FRR and bar charts of EERs. Fig. 5 shows the error curve calculated for the baseline system (facial surface depth maps) using the four distance measures described in section 6.

The results clearly show that diving by eigenvalues to normalise vector dimensions prior to calculating distance values significantly decreases error rates for both the euclidean and cosine distance measures, with the mahalanobis distance providing the lowest EER. The same four curves were produced for all surface representations described in section 4 and the EERs taken as a single comparative value, presented in Fig. 6.

It is clear from the EERs shown in Fig. 6, that surface gradient representations provide the most distinguishing information for face recognition. The horizontal derivatives give the lowest error rates of all, using the weighted cosine distance metric. In fact, the weighted cosine distance returns the lowest error rates for the majority of surface representations, except for a few cases when the weighted cosine EER is particularly high. However, which is the most effective surface representation seems to be dependent on the distance metric used for comparison (see Fig. 7), except for curvature representations, which are generally less distinguishing, regardless of the distance metric used.

Due to the orthogonal nature of the most effective surface representations (horizontal and vertical derivatives), we hypothesize that combing these representations will reduce error rates further. Therefore, in addition to the systems shown in Fig. 6, we test a number of system combinations by concatenating the face-keys projected from numerous surface spaces, attempting to utilise distinguishing features from multiple surface representations. The results for which are shown in Table 1, calculated by applying the weighted cosine distance measure to the extended face-keys combinations.

Table 1. Equal error rates of surface space combination systems

Surface Space Combinations	EER
Sobel X, Sobel Y, Horizontal gradient large, vertical gradient	12.1%
Laplacian, Horizontal gradient large, vertical gradient large	11.6%
Laplacian, Sobel X, Horizontal gradient, Horizontal gradient large, vertical gradient, vertical gradient large	11.4%

We have shown that a well-known two-dimensional face recognition method can be adapted for use on three-dimensional face models. Tests have been carried out on a large database of three-dimensional facial surfaces, captured under conditions that present typical difficulties when performing recognition. The error rates produced from baseline three-dimensional systems are significantly lower that those gathered in similar experiments using two-dimensional images. It is clear that three-dimensional face recognition has distinct advantages over conventional two-dimensional approaches.

Experimenting with a number of surface representations, we have discovered that facial surface gradient is more effective for recognition than depth and curvature representations. In particular, horizontal gradients produce the lowest error rates. This seems to indicate that horizontal derivatives provide more discriminatory information than vertical profiles. Another advantage is that gradients are likely to be more robust to inaccuracies in the alignment procedure, as the derivatives will be invariant to translations along the Z-axis.

Curvature representations do not seem to contain as much discriminatory information as the other surface representations. We find this surprising, as second derivatives should be less sensitive to inaccuracies of orientation and translation along the Z-axis. However, this could be a reflection

of inadequate 3D model resolution, which could be the cause of the noisy curvature images in Figure 8.

Testing three distance metrics has shown that the choice of method for face-key comparisons has a considerable affect on the resulting error rates. It is also evident that dividing each face-key by its respective eigenvalues, normalising dimensional distribution, usually improves results for both euclidean and cosine distances. This indicates that dimensional distribution is not necessarily proportional to discriminating ability and that surface space as a whole becomes more discriminative when distributed evenly. However, this is not the case for some of surface representations with higher EERs, suggesting that these representations incorporate only a few dominant useful components, which become masked when normalised with the majority of less discriminatory components.

The weighted cosine distance produces the lowest error rates for the majority of surface representations, including the optimum system. This metric has also provided the means to combine multiple face-keys, in an attempt to utilise advantages offered by numerous surfaces representations, reducing error rates further.

We have managed to reduce error rates from 17.8% EER, obtained using the initial depth maps, to an EER of 12.1% when the most effective surface representations were combined into a single system. These results are substantially lower than the best two-dimensional systems tested under similar circumstances in our previous investigations, proving that geometric face structure is useful for recognition when used independently from colour and texture information and capable of achieving high levels of accuracy. Given that the data capture method produces face models invariant to lighting conditions

and provides the ability to recognise faces regardless of pose, makes this system particularly attractive for use in security and surveillance applications.

In this specification, the verb "comprise" has its normal dictionary meaning, to denote non-exclusive inclusion. That is, use of the word "comprise" (or any of its derivatives) to include one feature or more, does not exclude the possibility of also including further features.

The reader's attention is directed to all and any priority documents identified in connection with this application and to all and any papers and documents which are filed concurrently with or previous to this specification in connection with this application and which are open to public inspection with this specification, and the contents of all such papers and documents are incorporated herein by reference.

All of the features disclosed in this specification (including any accompanying claims, abstract and drawings), and/or all of the steps of any method or process so disclosed, may be combined in any combination, except combinations where at least some of such features and/or steps are mutually exclusive.

Each feature disclosed in this specification (including any accompanying claims, abstract and drawings), may be replaced by alternative features serving the same, equivalent or similar purpose, unless expressly stated otherwise. Thus, unless expressly stated otherwise, each feature disclosed is one example only of a generic series of equivalent or similar features.

The invention is not restricted to the details of the foregoing embodiment(s). The invention extends to any novel one, or any novel

combination, of the features disclosed in this specification (including any accompanying claims, abstract and drawings), or to any novel one, or any novel combination, of the steps of any method or process so disclosed.

References

1. Turk, M., Pentland, A.: Eigenfaces for Recognition. Journal of Cognitive Neuroscience, Vol. 3, (1991) 72-86
2. Turk, M., Pentland, A.: Face Recognition Using Eigenfaces. In Proc. IEEE Conf. on Computer Vision and Pattern Recognition, (1991) 586-591
3. Beumier, C., Acheroy, M.: Automatic 3D Face Authentication. Image and Vision Computing, Vol. 18, No. 4, (2000) 315-321
4. Beumier, C., Acheroy, M.: Automatic Face Verification from 3D And Grey Level Clues. 11th Portuguese Conference on Pattern Recognition, 2000.
5. Gordon, G.: Face Recognition Based on Depth and Curvature Features. In Proc. of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, Champaign, Illinois (1992) 108-110.
6. Chua, C., Han, F., Ho, T.: 3D Human Face Recognition Using Point Signature. Proc. Fourth IEEE International Conference on Automatic Face and Gesture Recognition, (2000) 233-8
7. Zhao, W., Chellaa, R.: 3D Model Enhanced Face Recognition. In Proc. Of the International Conference on Image Processing, Vancouver (2000)
8. Romdhani, S., Blanz, V., Vetter, T.: Face Identification by Fitting a 3D Morphable Model using Linear Shape and Texture Error Functions. The European Conference on Computer Vision (2002)
9. Blanz, V., Romdhani, S., Vetter, T.: Face Identification across Different Poses and Illuminations with a 3D Morphable Model. In Proc. of the 5th IEEE Conference on AFGR (2002)
10. Phillips, P.J., Grother, P., Micheals, R.J., Blackburn, D.M., Tabassi, E., Bone, J.M.: FRVT 2002: Overview and Summary. http://www.frvt.org/FRVT2002/documents.htm, March (2003)
11. Phillips, P.J., Wechsler, H., Huang, J., Rauss, P.: The FERET database and evaluation procedure for face recognition algorithms Image and Vision Computing. J, Vol. 16, No. 5, (1998) 295-306
12. A. M. Coombes, R. Richards, A. Linney, V. Bruce, R. Fright. Description and recognition of faces from 3D data. Proc. of The International Society for Optical Engineering. Vol. 1766, (1992) 307-19.

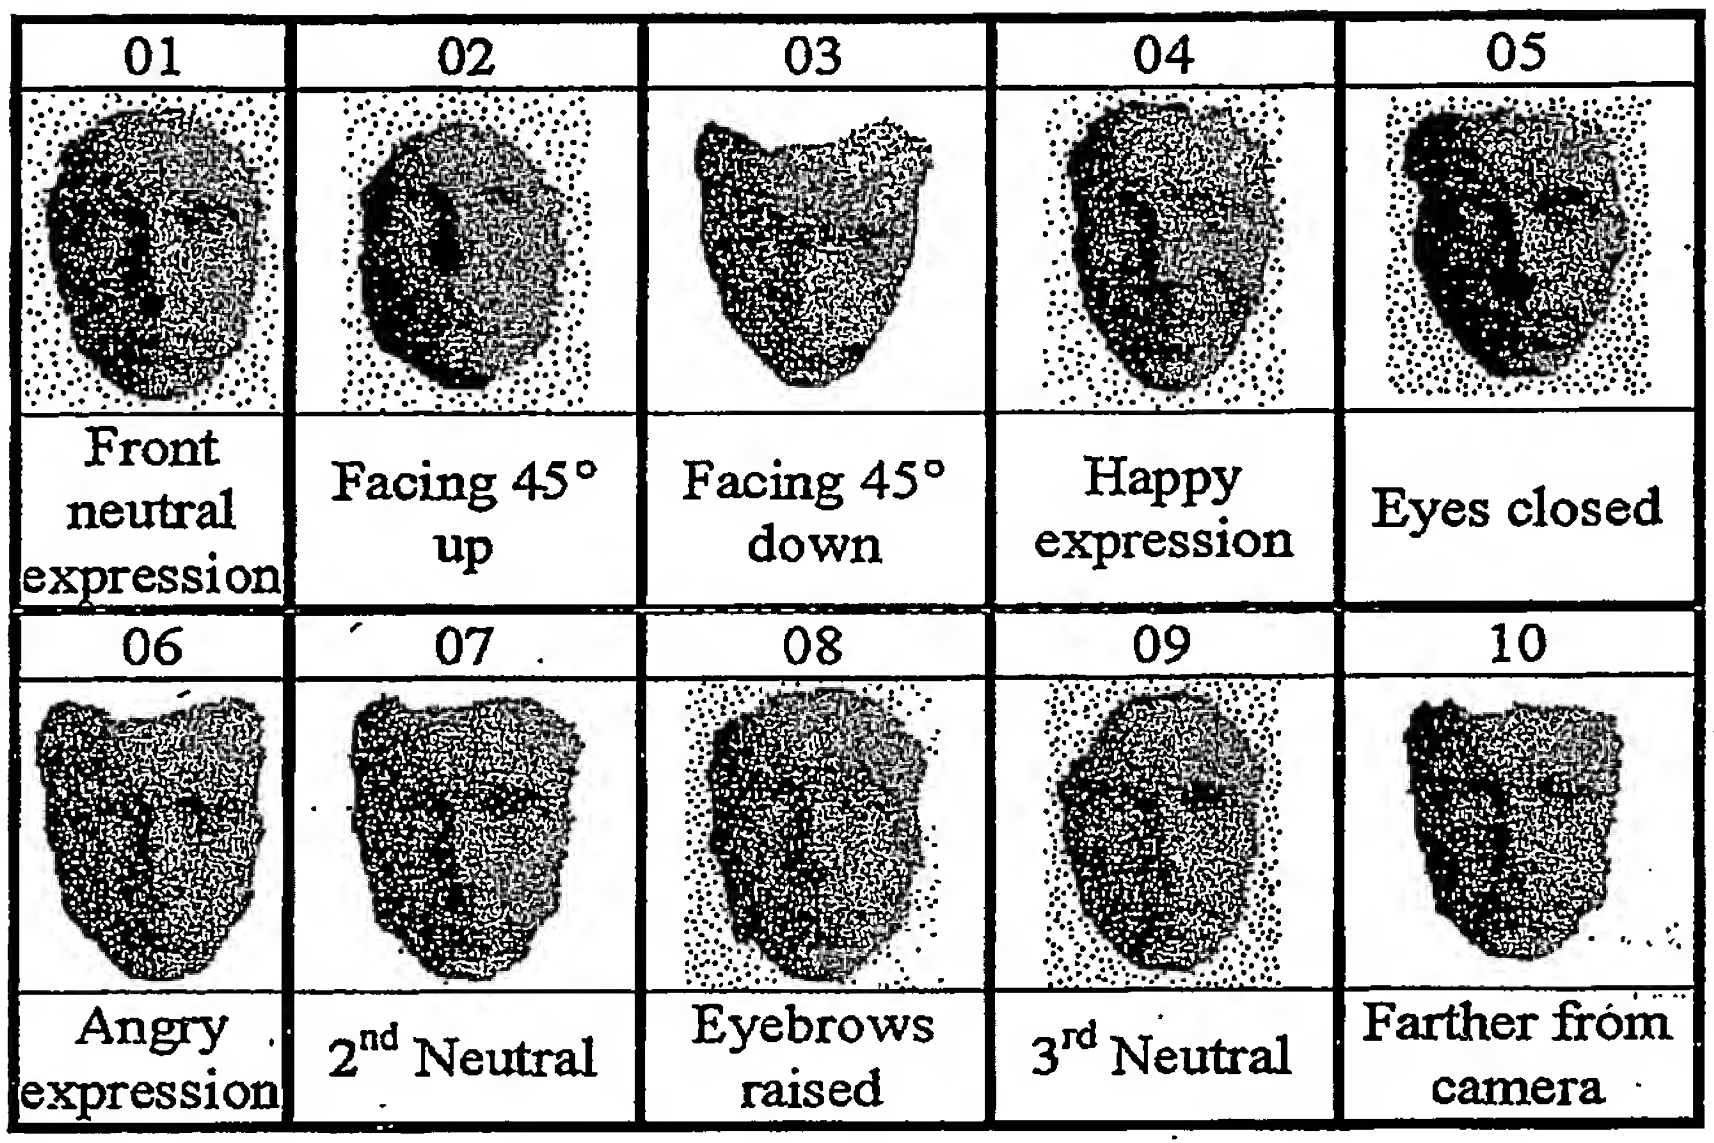

Fig. 1. Example face models taken from a 3D face database

Fig. 2. Orientation of a raw 3D face model (*left*) to a frontal pose (*middle*) and facial surface depth map (*right*)

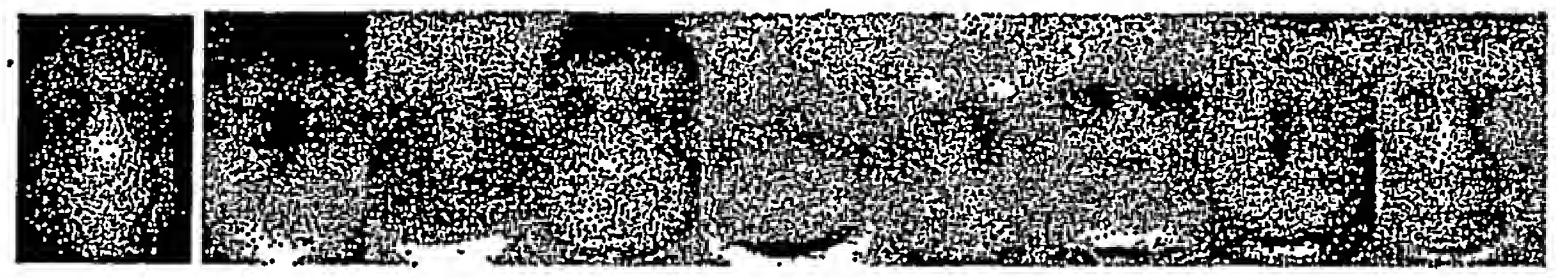

Fig. 3. Average depth map (*left most*) and first eight eigensurfaces

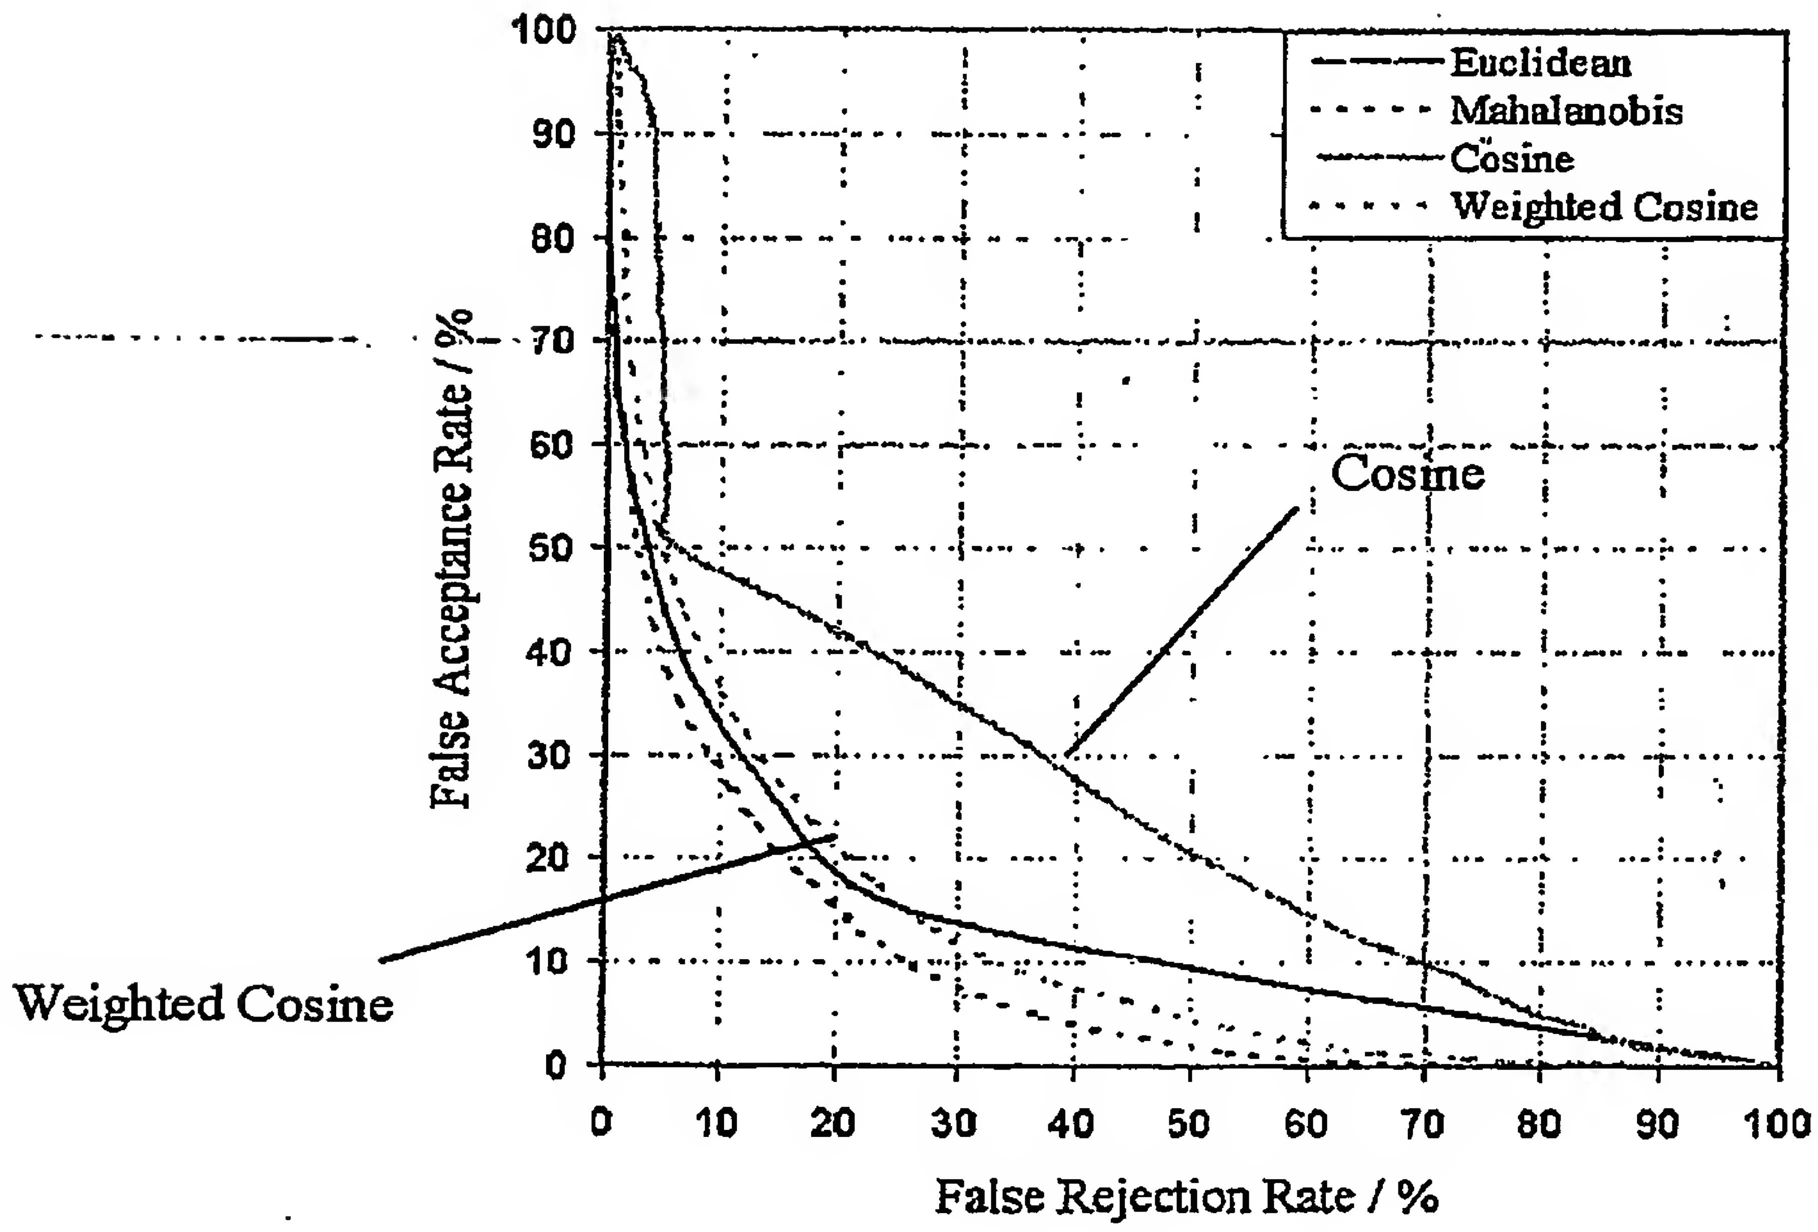

Fig. 4. Baseline 3D face recognition systems using facial surface depth maps and a range of distance metrics

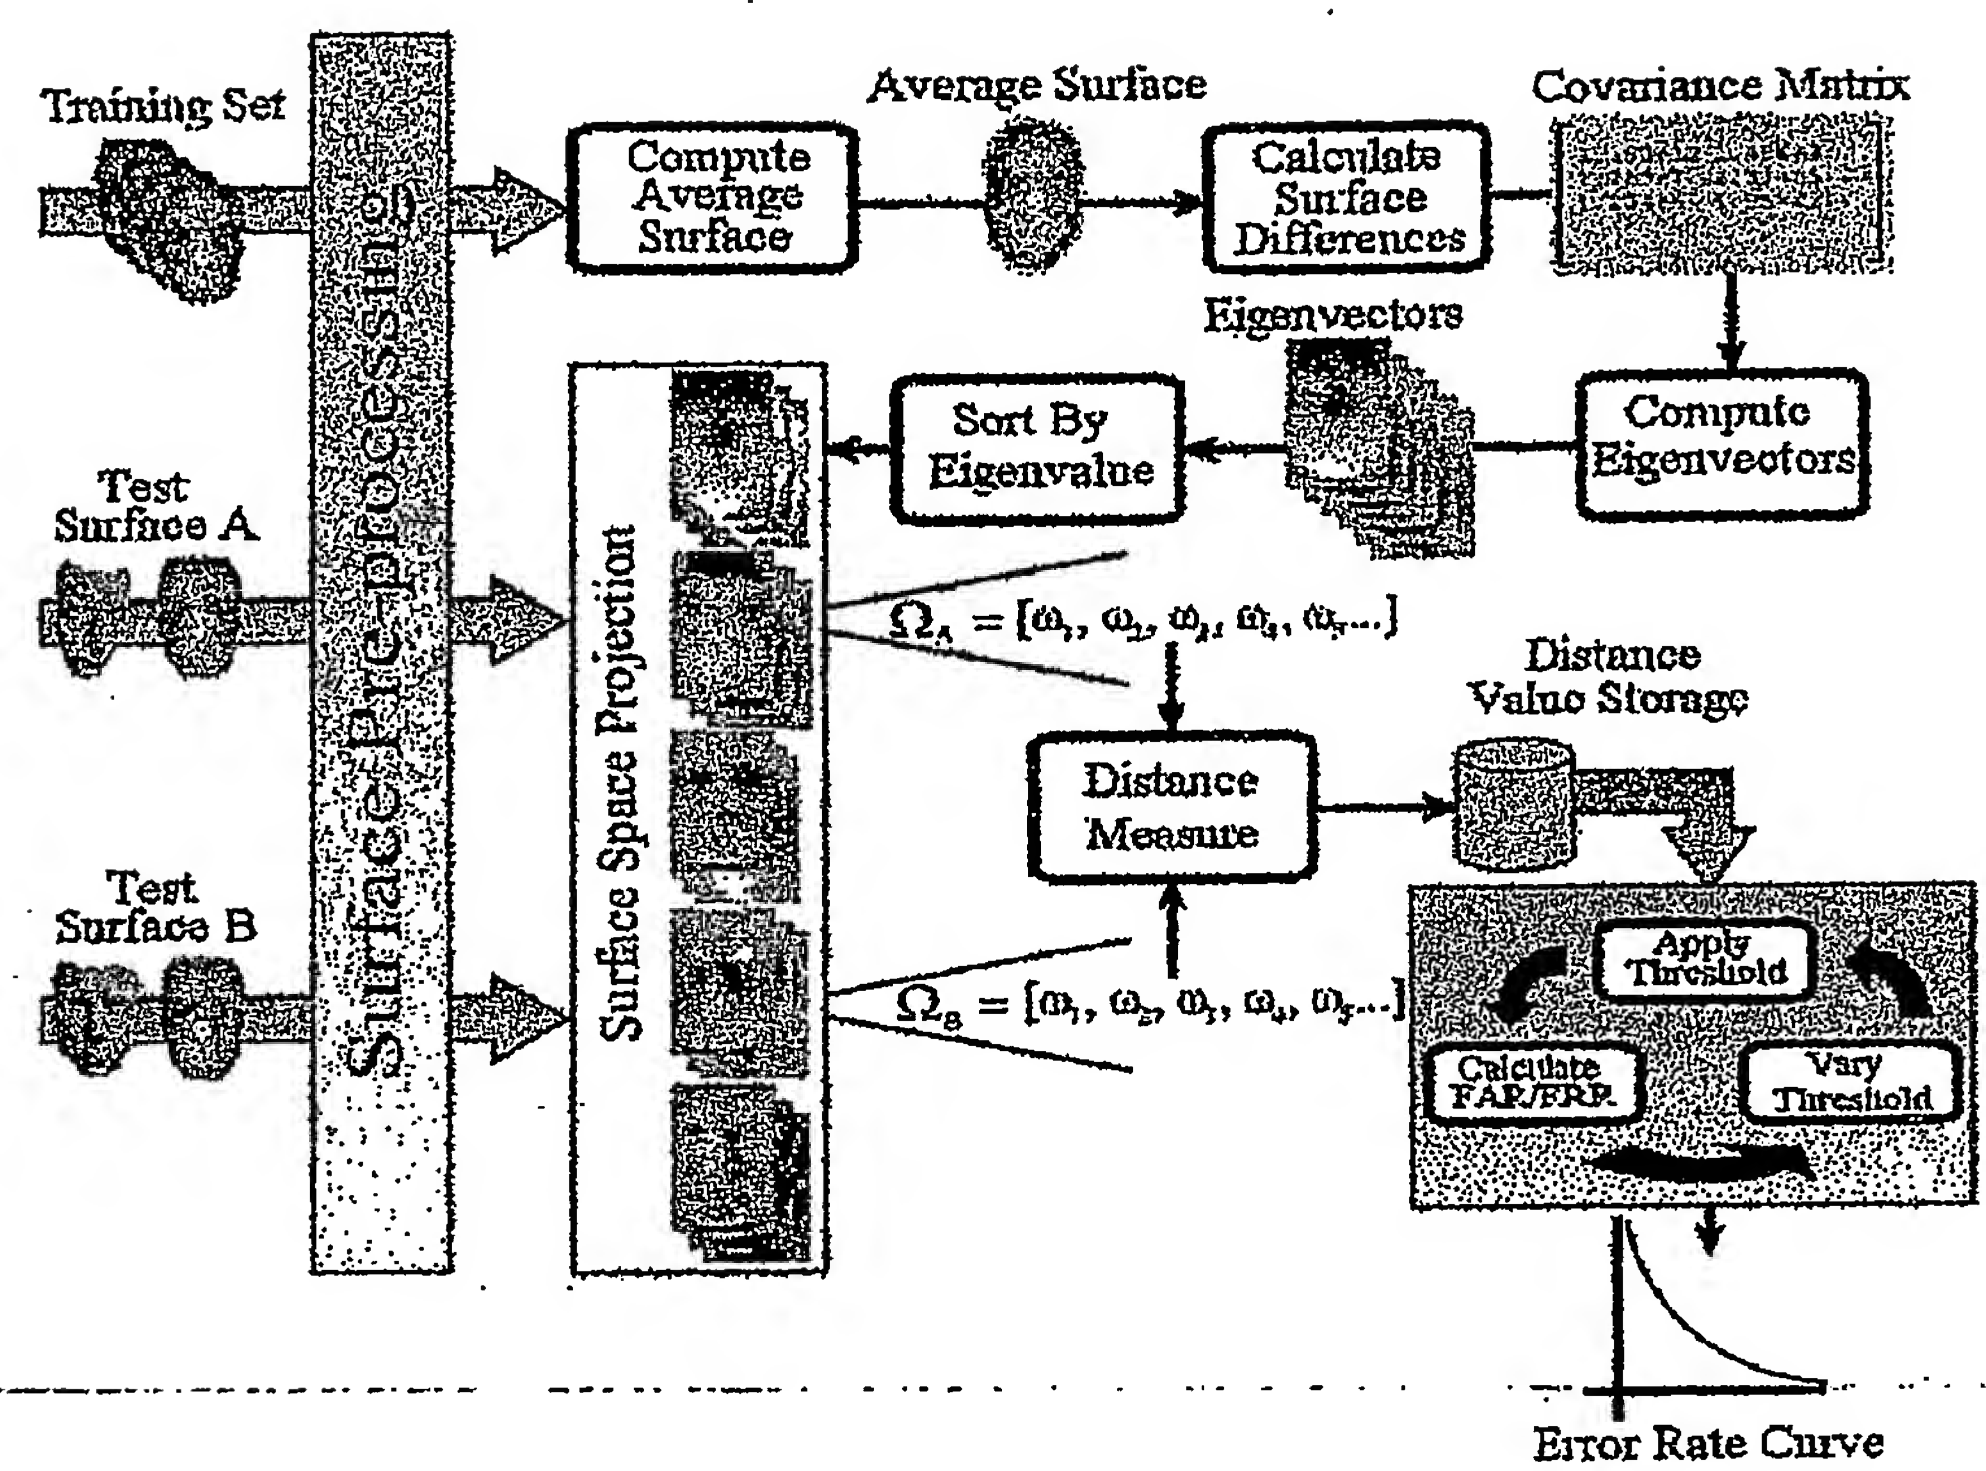

Fig. 5. Diagram of verification test procedure

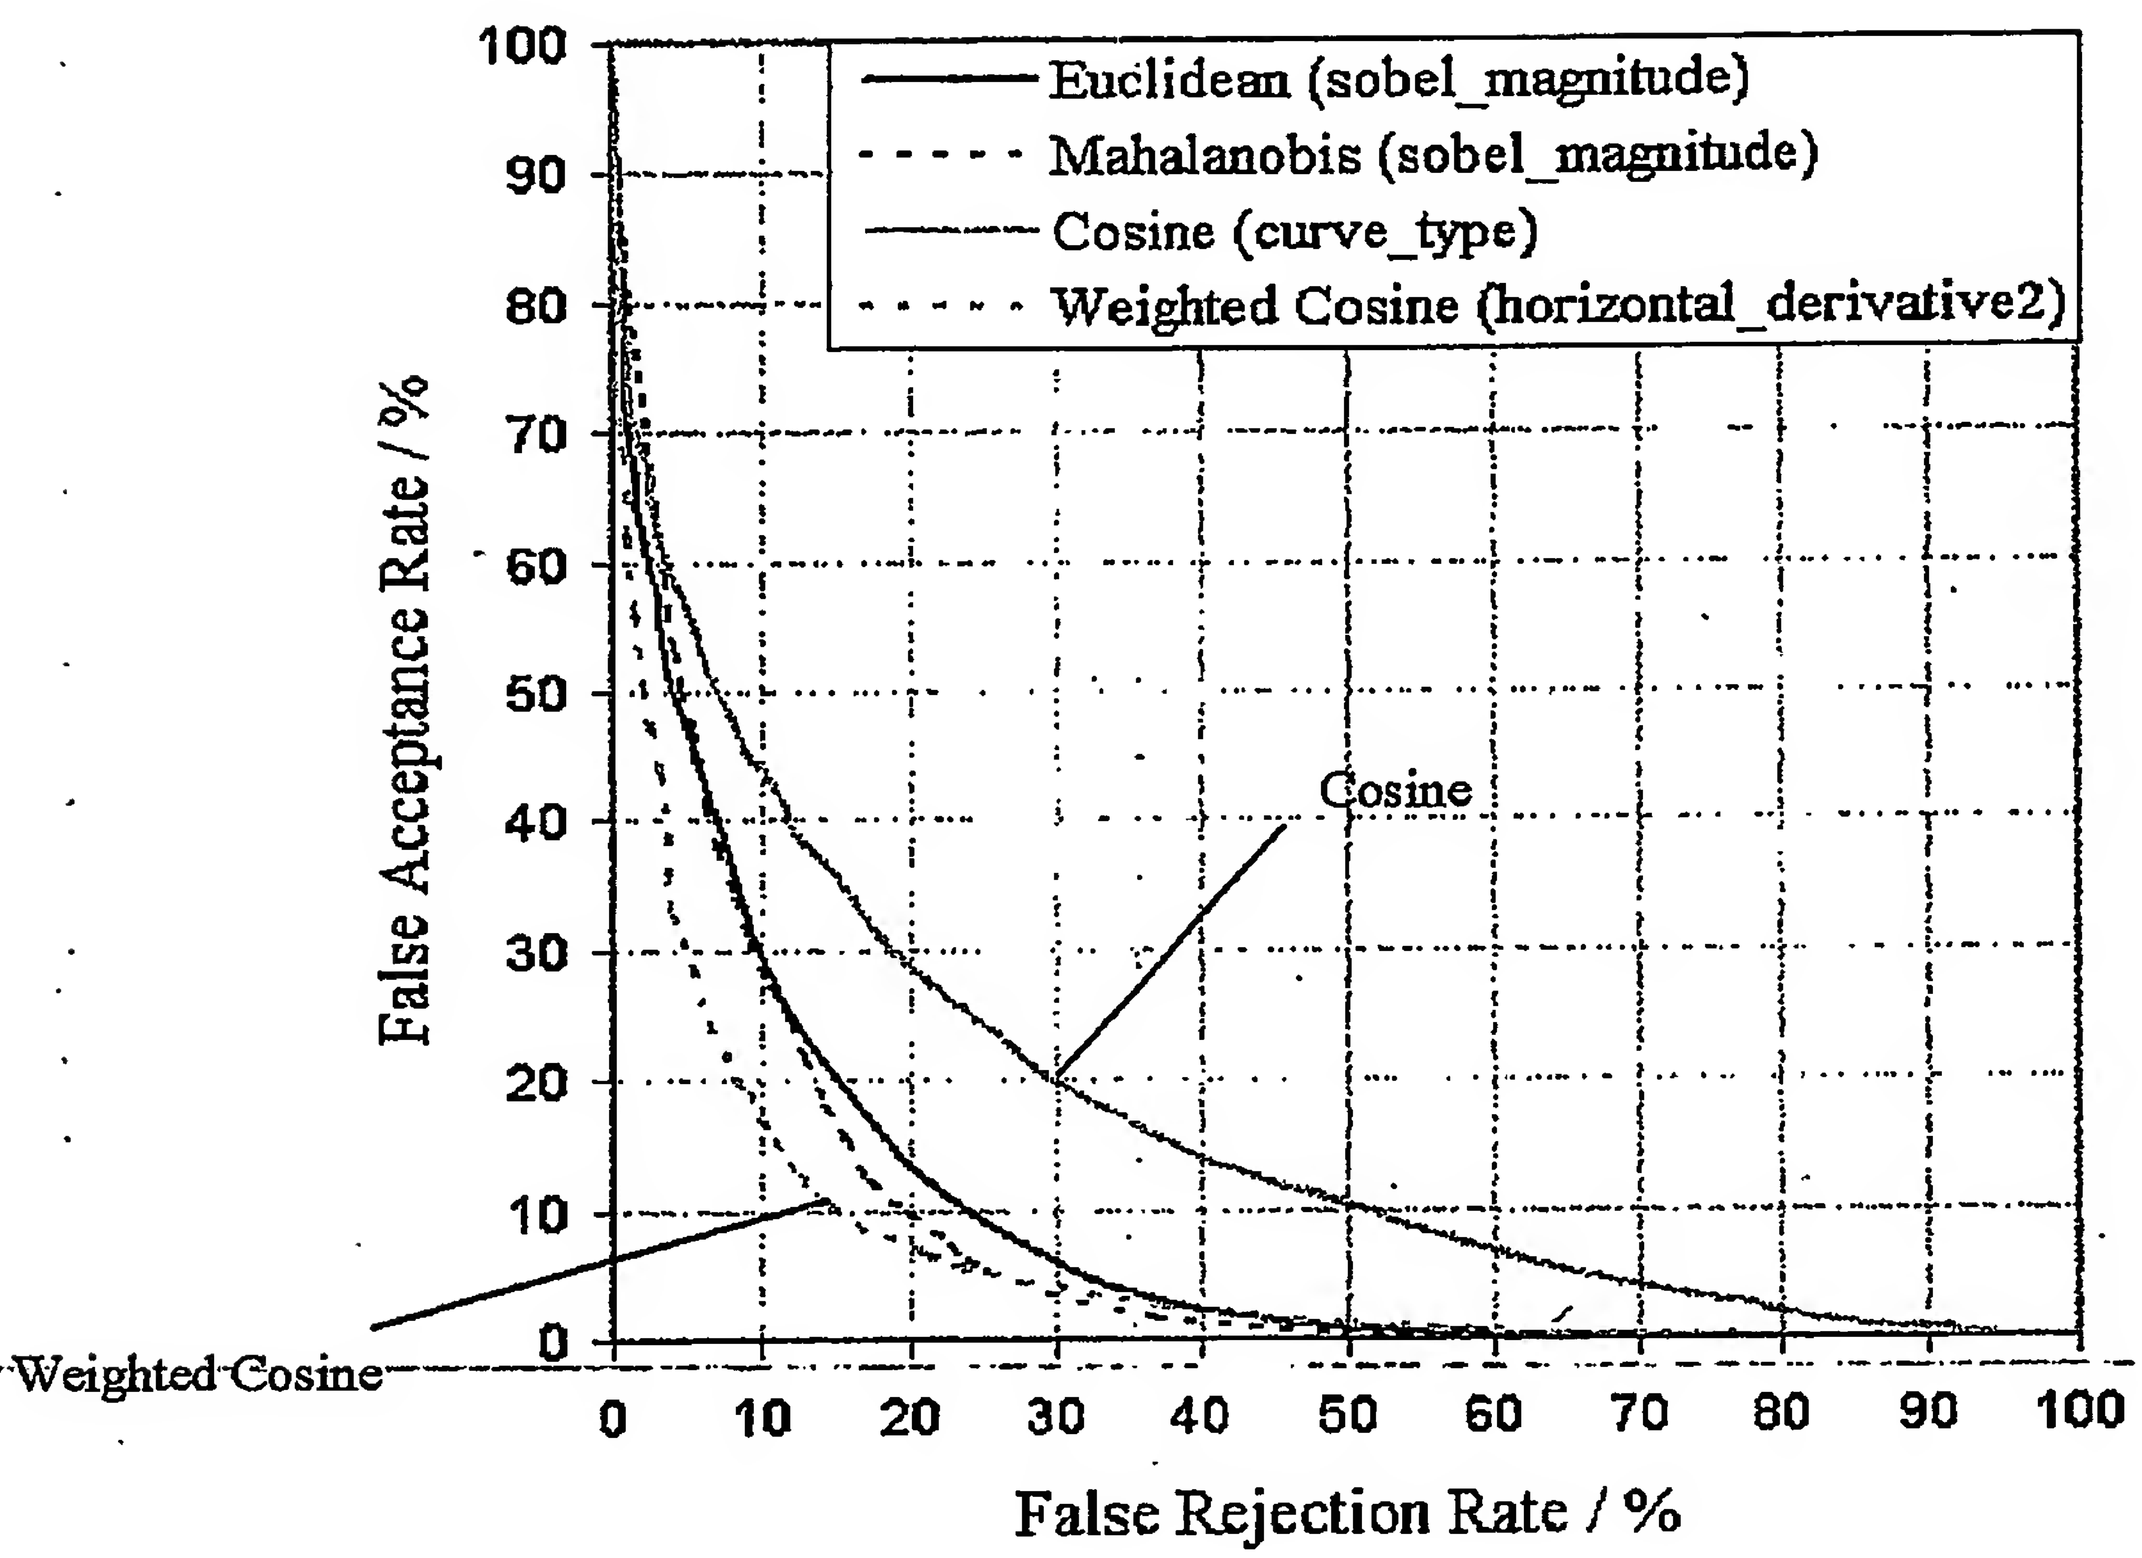

Fig. 6. Error rates of 3D face recognition systems using optimum surface representations and distance metrics.

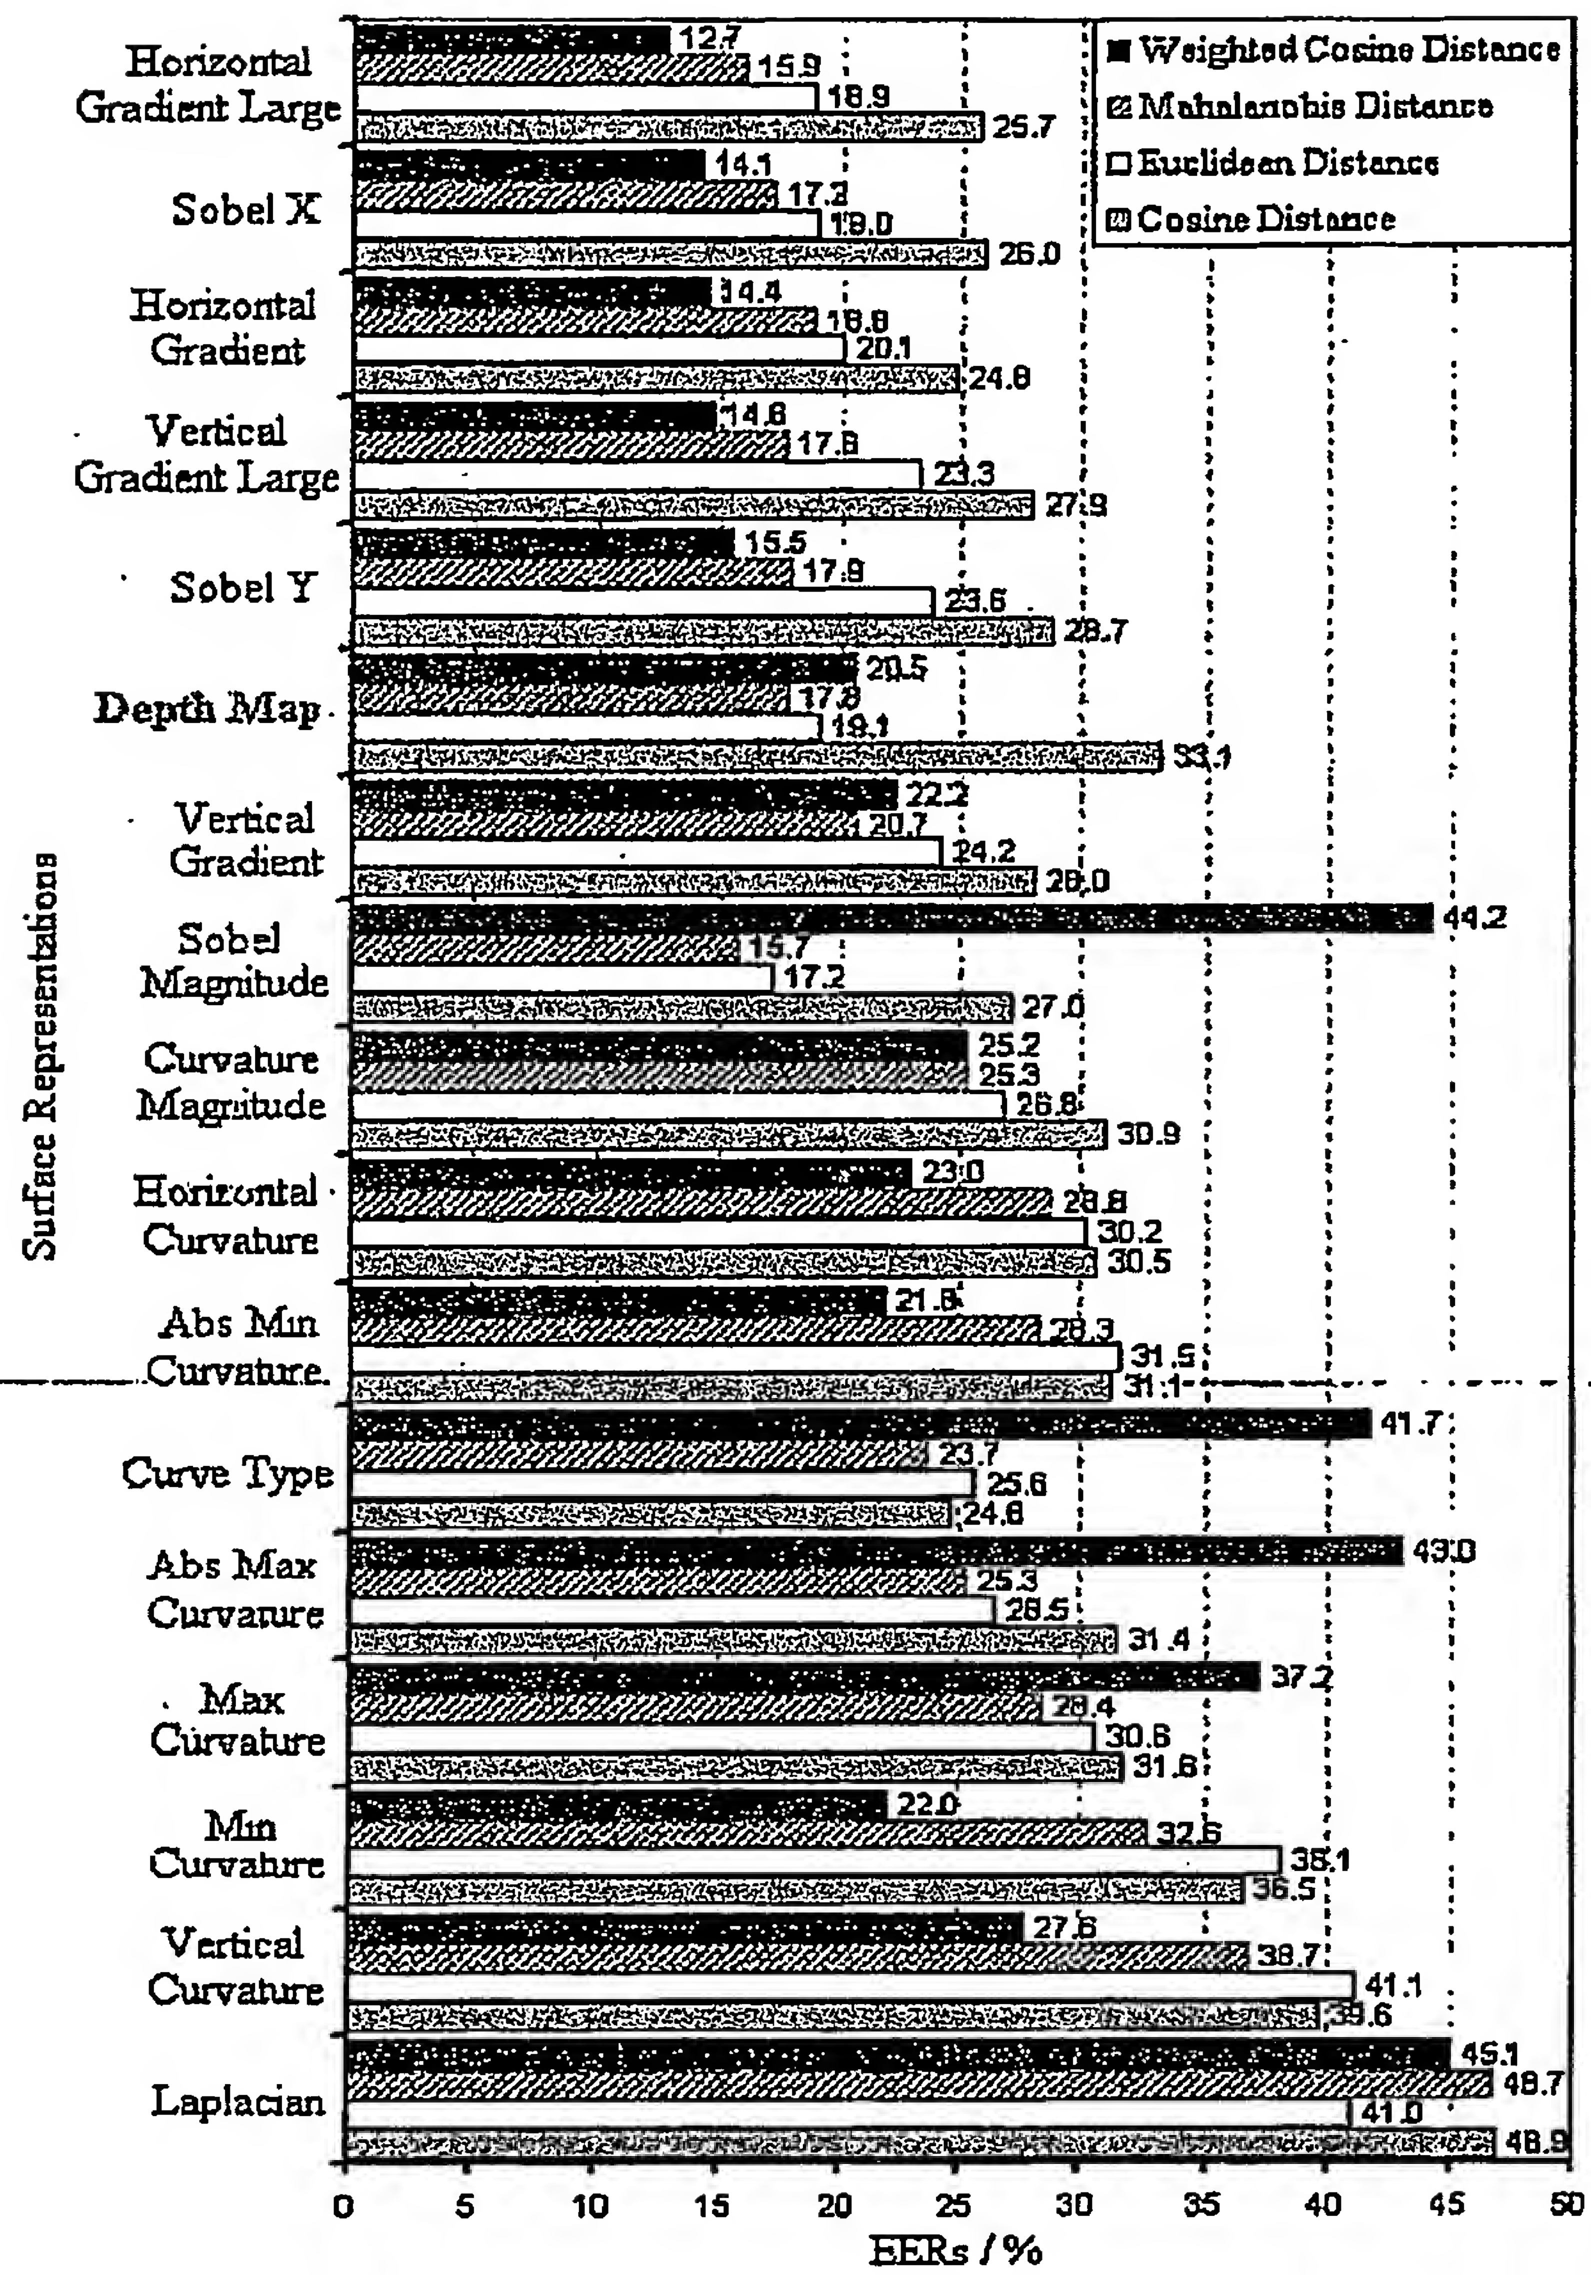

Fig. 7. Equal error rates of 3D face recognition systems using a variety of surface representations and distance metrics

Horizontal Gradient	Vertical Gradient	Horizontal Gradient Large	Vertical Gradient Large
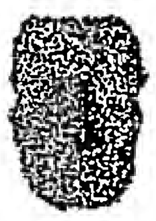 -1 1	 -1 1	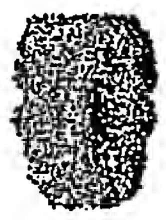 -1 0 0 0 1	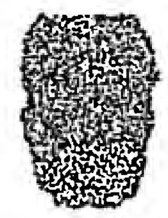 -1 0 0 0 1
Applies the 2x1 kernel to compute the horizontal derivative	Applies the 1x2 kernel to compute the vertical derivative	Horizontal gradient calculated over a greater horizontal distance	Vertical gradient calculated over a greater vertical distance
Laplacian	Sobel X	Sobel Y	Sobel Magnitude
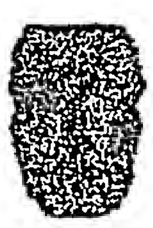 0 1 0 1 −4 1 0 1 0	 -1 0 1 -2 0 2 -1 0 1	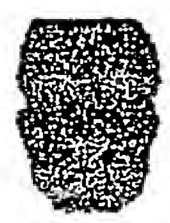 1 2 1 0 0 0 -1 -2 -1	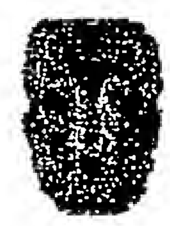
An isotropic measure of the second spatial derivative	Application of the sobel derivative filter in the horizontal direction	Application of the sobel derivative filter in the vertical direction	The magnitude of the X and Y sobel derivatives
Horizontal Curvature	Vertical Curvature	Curvature Magnitude	Curve Type
		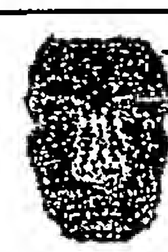	
Applies the sobel X kernel twice to calculate the second horizontal derivative	Applies the sobel Y kernel twice to calculate the second vertical derivative	The magnitude of the vertical and horizontal curvatures	Segmentation of the surface into 8 discreet curvature types
Min Curvature	Max Curvature	Abs Min Curvature	Abs Max Curvature
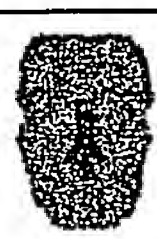	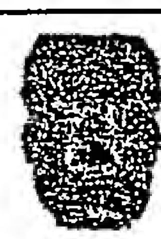	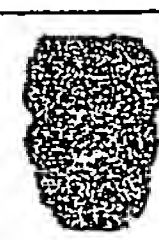	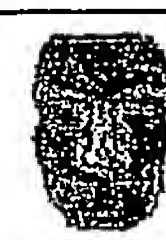
The minimum of the horizontal and vertical curvature values	The maximum of the horizontal and vertical curvature values	The minimum of the absolute horizontal and vertical curvatures	The maximum of the absolute horizontal and vertical curvatures

Figure 8. Brief descriptions of surface representations with the convolution kernels used.

PCT/EP2004/052502